Assignment: A1

Due: 6 September 2012

You are to explore the use of the Multi-Layer Perceptron (MLP), also called a neural net, to classify scanned images of the 26 lower-case characters (i.e., a-z). Several aspects of the MLP which deserve careful attention:

- **Input vector**: there are several possibilities, including:
  - Pixel values (binary): take the image pixels and make a 1D vector
  - Low-level features: e.g., area
  - High-level features: e.g., curvature of the boundary
- **Hidden layers**:
  - How many hidden layers
  - How many nodes per layer
- **Algorithm parameters**:
  - Stopping criterion
  - Weights on gradient descent
- **Data Management**:
  - How to select training, testing, and validation sets
  - How to use cross-validation

In addition, the results need to be presented in a strong statistical framework; this means computing statistics (e.g., mean, variance) over several trials (how many?), and showing confidence intervals.

Finally, the analysis and interpretation are the essential parts of the report; use these to present your findings, understanding and remaining problems.

There is a set of sample images on the class data sub-directory.